

EL685270603

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**System and Method Providing Improved
Head Motion Estimations for Animation**

Inventor(s):
Zicheng Liu
Zhengyou Zhang

ATTORNEY'S DOCKET NO. MS1-684US

09344130600

TECHNICAL FIELD

The disclosure below relates to the recovery of face shape from images and the facial animation. More particularly, the following description relates to determining the head motion between two images based on symmetrical features in the two images.

BACKGROUND

One of the most interesting and difficult problems in computer graphics is the effortless generation of realistic looking, animated human face models. Animated face models are essential to computer games, film making, online chat, virtual presence, video conferencing, etc. So far, the most popular commercially available tools have utilized laser scanners. Not only are these scanners expensive, the data are usually quite noisy, requiring hand touchup and manual registration prior to animating the model. Because inexpensive computers and cameras are widely available, there is a great interest in producing face models directly from images. In spite of progress toward this goal, the available techniques are either manually intensive or computationally expensive.

Facial modeling and animation has been a computer graphics research topic for over 25 years [6, 16, 17, 18, 19, 20, 21, 22, 23, 27, 30, 31, 33]. The reader is referred to Parke and Waters' book [23] for a complete overview.

Lee et al. [17, 18] developed techniques to clean up and register data generated from laser scanners. The obtained model is then animated using a physically based approach.

DeCarlo et al. [5] proposed a method to generate face models based on face measurements randomly generated according to anthropometric statistics. They

1 showed that they were able to generate a variety of face geometries using these
2 face measurements as constraints.

3 A number of researchers have proposed to create face models from two
4 views [1, 13, 4]. They all require two cameras which must be carefully set up so
5 that their directions are orthogonal. Zheng [37] developed a system to construct
6 geometrical object models from image contours, but it requires a turn-table setup.

7 Pighin et al. [26] developed a system to allow a user to manually specify
8 correspondences across multiple images, and use vision techniques to computer 3D
9 reconstructions. A 3D mesh model is then fit to the reconstructed 3D points. They
10 were able to generate highly realistic face models, but with a manually intensive
11 procedure.

12 Blanz and Vetter [3] demonstrated that linear classes of face geometries and
13 images are very powerful in generating convincing 3D human face models from
14 images. Blanz and Vetter used a large image database to cover every skin type.

15 Kang et al. [14] also use linear spaces of geometrical models to construct
16 3D face models from multiple images. But their approach requires manually
17 aligning the generic mesh to one of the images, which is in general a tedious task
18 for an average user.

19 Fua et al. [8] deform a generic face model to fit dense stereo data, but their
20 face model contains a lot more parameters to estimate because basically all of the
21 vertexes are independent parameters, plus reliable dense stereo data are in general
22 difficult to obtain with a single camera. Their method usually takes 30 minutes to
23 an hour, while ours takes 2-3 minutes.

24 Guenter et al. [9] developed a facial animation capturing system to capture
25 both the 3D geometry and texture image of each frame and reproduce high quality

facial animations. The problem they solved is different from what is addressed here in that they assumed the person's 3D model was available and the goal was to track the subsequent facial deformations.

SUMMARY

The system described below provides improved procedures to estimate head motion between two images of a face. A procedure is described that first, identifies locations of a number of distinct facial features in two images. The procedure of estimating head motion with respect to these locations corresponds to the determination of a number of unknown quantities. For example, these identified locations can correspond to the eye corners, mouth corners and nose tip.

Next, the locations are converted into as a set of physical face parameters based on the symmetry of the identified distinct facial features. The set of physical parameters reduces the number of unknowns as compared to the number of equations used to determine the unknowns. This reduction in unknowns and relative increase in the number of equations used to determine the unknowns increases the redundancy and thus, the robustness of the head motion estimation.

Finally, the points corresponding to the identified points between the two images are used to solve the head motion. The head motion estimation operation includes: (a) estimating each of the set of physical parameters, (b) estimating a first head pose transform corresponding to the first image, and (c) estimating a second head pose transform corresponding to the second image. One of the physical parameters is set to a constant due to the fact that a scale cannot be determined from the two images.

1 Optionally, an inequality constraint can be placed on a particular physical
2 parameter, such that the parameter is constrained within a predetermined minimum
3 and maximum value. For example, one description provides that the parameter
4 corresponds to a nose tip. The predetermined minimum value is zero (0) and the
5 predetermined maximum value is a reasonable value based on absolute values of
6 other locations – recall that each location identifies a facial feature. The inequality
7 constraint is converted to an equality constraint by using a penalty function. Then,
8 the inequality constraint is used during the head motion estimation to add
9 additional robustness to the motion estimation.

10 In yet another description, a procedure is provided to use the head motion
11 estimation described above as an initial estimation in combination with a feature
12 matching algorithm. To accomplish this, the procedure first involves identifying
13 locations of a plurality of distinct facial features in the two images, the locations
14 corresponding to a number of unknowns determined upon estimation of head
15 motion. For example, these identified locations can correspond to the eye corners,
16 mouth corners and nose tip.

17 Next, the identified locations are converted into a set of physical face
18 parameters based on the symmetry of the identified distinct facial features, the set
19 of physical parameters reducing the number of unknowns. Next, a first set of
20 matched points (corresponding to previously identified distinct facial features) is
21 used to determine the head motion and the physical facial parameters. Finally, a
22 second set of matched points from a feature matching algorithm is incorporated,
23 together with the first set of matched points, to refine the estimation of the head
24 motion and the physical facial parameters. These novel procedures provide
25

1 additional robustness to systems that estimate head motion during animation
2 modeling.

3 4 **BRIEF DESCRIPTION OF THE DRAWINGS**

5 Fig. 1 is a block diagram of a computer system capable of performing the
6 operations described below.

7 Fig. 2 illustrates how to mark facial features on an image.

8 Figs. 3, 5, 6 are flow charts showing sequences of actions for creating a 3D
9 face model.

10 Fig. 4 shows the selection of different head regions as described below.

11 Fig. 7 illustrates a coordinate system that is based on symmetry between
12 selected feature points on an image.

13 14 **DETAILED DESCRIPTION**

15 The following description sets forth a specific embodiment of a 3D
16 modeling system that incorporates elements recited in the appended claims. The
17 embodiment is described with specificity in order to meet statutory requirements.
18 However, the description itself is not intended to limit the scope of this patent.
19 Rather, the claimed invention might eventually be embodied in other ways, to
20 include different elements or combinations of elements similar to the ones
21 described in this document, in conjunction with other present or future
22 technologies.

System Overview

Fig. 1 shows components of our system. The equipment includes a computer 10 and a video camera 12. The computer is a typical desktop, laptop, or similar computer having various typical components such as a keyboard/mouse, display, processor, peripherals, and computer-readable media on which an operating system and application programs are stored and from which the operating system and application programs are executed. Such computer-readable media might include removable storage media, such as floppy disks, CDROMs, tape storage media, etc. The application programs in this example include a graphics program designed to perform the various techniques and actions described below.

The video camera is an inexpensive model such as many that are widely available for Internet videoconferencing. We assume the intrinsic camera parameters have been calibrated, a reasonable assumption given the simplicity of calibration procedures [36].

Data Capture

The first stage is data capture. The user takes two images with a small relative head motion, and two video sequences: one with the head turning to each side. Alternatively, the user can simply turn his/her head from left all the way to the right, or vice versa. In that case, the user needs to select one approximately frontal view while the system automatically selects the second image and divides the video into two sequences. In the sequel, we call the two images the *base images*.

The user then locates five markers in each of the two base images. As shown in Fig. 2, the five markers correspond to the two inner eye corners 20, nose tip 21, and two mouth corners 22.

The next processing stage computes the face mesh geometry and the head pose with respect to the camera frame using the two base images and markers as input.

The final stage determines the head motions in the video sequences, and blends the images to generate a facial texture map.

Notation

We denote the homogeneous coordinates of a vector \mathbf{x} by $\tilde{\mathbf{x}}$, i.e., the homogeneous coordinates of an image point $\mathbf{m} = (u, v)^T$ are $\tilde{\mathbf{m}} = (u, v, 1)^T$, and those of a 3D point $\mathbf{p} = (x, y, z)^T$ are $\tilde{\mathbf{p}} = (x, y, z, 1)^T$. A camera is described by a pinhole model, and a 3D point \mathbf{p} and its image point \mathbf{m} are related by

$$\lambda \tilde{\mathbf{m}} = \mathbf{A} \mathbf{P} \Omega \tilde{\mathbf{p}}$$

where λ is a scale, and \mathbf{A} , \mathbf{P} , and Ω are given by

$$\mathbf{A} = \begin{pmatrix} \alpha & \lambda & u_0 \\ 0 & \beta & v_0 \\ 0 & 0 & 1 \end{pmatrix} \quad \mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad \Omega = \begin{pmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{pmatrix}$$

The elements of matrix \mathbf{A} are the intrinsic parameters of the camera and matrix \mathbf{A} maps the normalized image coordinates to the pixel image coordinates (see e.g. [7]). Matrix \mathbf{P} is the perspective projection matrix. Matrix Ω is the 3D rigid transformation (rotation \mathbf{R} and translation \mathbf{t}) from the object/world coordinate system to the camera coordinate system. When two images are concerned, a prime ' is added to denote the quantities related to the second image.

The fundamental geometric constraint between two images is known as the *epipolar constraint* [7, 35]. It states that in order for a point \mathbf{m} in one image and a point \mathbf{m}' in the other image to be the projections of a single physical point in space, or in other words, in order for them to be matched, they must satisfy

$$\tilde{\mathbf{m}}'^T \mathbf{A}'^{-T} \mathbf{E} \mathbf{A}^{-1} \tilde{\mathbf{m}} = 0$$

where $\mathbf{E} = [\mathbf{t}_r]_{\times} \mathbf{R}_r$ is known as the essential matrix, $(\mathbf{R}_r, \mathbf{t}_r)$ is the relative motion between the two images, and $[\mathbf{t}_r]_{\times}$ is a skew symmetric matrix such that $\mathbf{t}_r \times \mathbf{v} = [\mathbf{t}_r]_{\times} \mathbf{v}$ for any 3D vector \mathbf{v} .

Linear Class of Face Geometries

Instead of representing a face as a linear combination of real faces or face models, we represent it as a linear combination of a neutral face model and some number of face *metrics*, where a metric is a deformation vector that linearly deforms a face in a certain way, such as to make the head wider, make the nose bigger, etc. Each deformation vector specifies a plurality of displacements corresponding respectively to the plurality of 3D points of the neutral face model.

To be more precise, let's denote the face geometry by a vector $\mathbf{S} = (\mathbf{v}_1^T, \dots, \mathbf{v}_n^T)^T$, where $\mathbf{v}_i = (X_i, Y_i, Z_i)^T$, ($i=1, \dots, n$) are the vertices, and a metric by a vector $\mathbf{M} = (\delta \mathbf{v}_1, \dots, \delta \mathbf{v}_n)^T$, where $\delta \mathbf{v}_i = (\delta X_i, \delta Y_i, \delta Z_i)^T$. Given a neutral face $\mathbf{S}^0 = (\mathbf{v}_1^{0T}, \dots, \mathbf{v}_n^{0T})^T$, and a set of m metrics $\mathbf{M}^j = (\delta \mathbf{v}_1^{jT}, \dots, \delta \mathbf{v}_n^{jT})^T$, the linear space of face geometries spanned by these metrics is

$$\mathbf{S} = \mathbf{S}^0 + \sum_{j=1}^m c_j \mathbf{M}^j \text{ subject to } c_j \in [l_j, u_j]$$

where c_j 's are the metric coefficients and l_j and u_j are the valid range of c_j . In our implementation, the neutral face and all the metrics are designed by an

1 artist, and it is done only once. The neutral face contains 194 vertices and 360
2 triangles. There are 65 metrics.

4 **Image Matching and 3D Reconstruction**

5 We now describe our techniques to determine the face geometry from just
6 two views. The two base images are taken in a normal room by a static camera
7 while the head is moving in front. There is no control on the head motion, and the
8 motion is unknown. We have to determine first the motion of the head and match
9 some pixels across the two views before we can fit an animated face model to the
10 images. However, some preprocessing of the images is necessary.

12 **Determining Facial Portions of the Images**

13 Fig. 3 shows actions performed to distinguish a face in the two selected
14 images from other portions of the images.

15 There are at least three major groups of objects undergoing different
16 motions between the two views: background, head, and other parts of the body
17 such as the shoulder. If we do not separate them, there is no way to determine a
18 meaningful head motion, since the camera is static, we can expect to remove the
19 background by subtracting one image from the other. However, as the face color
20 changes smoothly, a portion of the face may be marked as background. Another
21 problem with the image subtraction technique is that the moving body and the head
22 cannot be distinguished.

23 An initial step 100 comprises using image subtraction to create the first
24 mask image in which pixels having different colors in the two base images are
25 marked.

1 A step 101 comprises identifying locations of a plurality of distinct facial
2 features in the base images. In this example, the user does this manually, by
3 marking the eyes, nose, and mouth, as described above and shown in Fig. 2.
4 Automated techniques could also be used to identify these points.

5 A step 102 comprises calculating a range of skin colors by sampling the
6 base images at the predicted portions, or at locations that are specified relative to
7 the user-indicated locations of the facial features. This allows us to build a color
8 model of the face skin. We select pixels below the eyes and above the mouth, and
9 computer a Gaussian distribution of their colors in the RGB space. If the color of
10 a pixel matches this face skin color model, the pixel is marked as a part of the face.

11 A step 103 comprises creating a second mask image that marks any image
12 pixels having colors corresponding to the calculated one or more skin colors.

13 Either union or intersection of the two mask images is not enough to locate
14 the face because it will include either too many (e.g., including undesired moving
15 body) or too few (e.g., missing desired eyes and mouth) pixels. Since we already
16 have information about the position of eye corners and mouth corners, we initially
17 predict the approximate boundaries of the facial portion of each image, based on
18 the locations identified by the user. More specifically, step 104 comprises
19 predicting an inner area and an outer area of the image. The outer area
20 corresponds roughly to the position of the subject's head in the image, while the
21 inner area corresponds roughly to the facial portion of the head.

22 Fig. 4 shows these areas, which are defined as ellipses. The inner ellipse 23
23 covers most of the face, while the outer ellipse 24 is usually large enough to
24 enclose the whole head. Let d_e be the image distance between the two inner eye
25 corners, and d_{em} , the vertical distance between the eyes and the mouth. The width

and height of the inner ellipse are set to $5d_e$ and $3d_{em}$. The outer ellipse is 25% larger than the inner one.

In addition, step 104 includes predicting or defining a lower area of the image that corresponds to a chin portion of the head. The lower area aims at removing the moving body, and is defined to be $0.6d_{em}$ below the mouth.

Within the inner ellipse, a “union” or “joining” operation 105 is used: we note all marked pixels in the first mask image and also any unmarked pixels of the first mask image that correspond in location to marked pixels in the second mask image. Between the inner and outer ellipses (except for the lower region), an image subtraction operation 106 is used: we note which pixels (marked or unmarked) of the image have different colors relative to correspondingly located pixels in the other image. In the lower part, we use an “intersection” operation 107: we note any marked pixels in the first mask image that correspond in location to marked pixels in the second mask image.

A step 108 comprises forming a final mask image that marks the noted pixels as being part of the head. This involves joining the mask image to the mask image of the other image. More specifically, the corresponding noted pixels of each base image are logically OR's to create the final mask image.

Corner Matching and Motion Determination

One popular technique of image registration is optical flow [12, 2], which is based on the assumption that the intensity/color is conserved. This is not the case in our situation: the color of the same physical point appears to be different in images because the illumination changes when the head is moving. We therefore resort to a feature-based approach that is more robust to intensity/color variations.

1 It consists of the following steps: (i) detecting corners in each image; (ii) matching
2 corners between the two images; (iii) detecting false matches based on a robust
3 estimation technique; (iv) determining the head motion; (v) reconstructing matched
4 points in 3D space.

5 Fig. 5 shows the sequence of operations.

6 **Corner Detection.** In a step 110, we use the Plessey corner detector, a
7 well-known technique in computer vision [10]. It locates corners corresponding to
8 high curvature points in the intensity surface if we view an image as a 3D surface
9 with the third dimension being the intensity. Only corners whose pixels are white
10 in the mask image are considered.

11 **Corner Matching.** In a step 111, for each corner in the first image we
12 choose an 11×11 window centered on it, and compare the window with windows
13 of the same size, centered on the corners in the second image. A zero-mean
14 normalized cross correlation between two windows is computed [7]. If we
15 rearrange the pixels in each window as a vector, the correlation score is equivalent
16 to the cosine angle between two intensity vectors. It ranges from -1 , for two
17 windows which are not similar at all, to 1 , for two windows which are identical. If
18 the largest correlation score exceeds a prefixed threshold (0.866 in our case), then
19 that corner in the second image is considered to be the *match candidate* of the
20 corner in the first image. The match candidate is retained as a *match* if and only if
21 its match candidate in the first image happens to be *the corner being considered*.
22 This symmetric test reduces many potential matching errors.

23 **False Match Detection.** Operation 112 comprises detecting and discarding
24 false matches. The set of matches established so far usually contains false matches
25 because correlation is only a heuristic. The only geometric constraint between two

1 images is the epipolar constraint $\tilde{\mathbf{m}}'^T \mathbf{A}'^{-T} \mathbf{E} \mathbf{A}^{-1} \tilde{\mathbf{m}} = 0$. If two points are correctly
2 matched, they must satisfy this constraint, which is unknown in our case.
3 Inaccurate location of corners because of intensity variation or lack of strong
4 texture features is another source of error. In a step 109, we use the technique
5 described in [35] to detect both false matches and poorly located corners and
6 simultaneously estimate the epipolar geometry (in terms of the essential matrix \mathbf{E}).
7 That technique is based on a robust estimation technique known as the *least*
8 *median squares* [28], which searches in the parameter space to find the parameters
9 yielding the smallest value for the *median* of squared residuals computed for the
10 entire data set. Consequently, it is able to detect false matches in as many as
11 49.9% of the whole set of matches.

12 13 Motion Estimation

14 In a step 113, we compute an initial estimate of the relative head motion
15 between two images, denoted by rotation \mathbf{R}_r and translation \mathbf{t}_r . If the image
16 locations of the identified feature points are precise, one could use a five-point
17 algorithm to compute camera motion from Matrix \mathbf{E} [7, 34]. Motion $(\mathbf{R}_r, \mathbf{t}_r)$ is then
18 re-estimated with a nonlinear least-squares technique using all remaining matches
19 after having discarded the false matches [34].

20 However, the image locations of the feature point are not usually precise. A
21 human typically cannot mark the feature points with high precision. An automatic
22 facial feature detection algorithm may not produce perfect results. When there are
23 errors, a five-point algorithm is not robust even when refined with a well-known
24 bundle adjustment technique.
25

For each of the five feature points, its 3D coordinates (x, y, z) coordinates need to be determined – fifteen (15) unknowns. Then, motion vector ($\mathbf{R}_t, \mathbf{t}_t$) needs to be determined – adding six (6) more unknowns. One unknown quantity is the magnitude, or global scale, which will never be determined from images alone. Thus, the number of unknown quantities that needs to be determined is twenty (i.e., $15+6-1 = 20$). The calculation of so many unknowns further reduces the robustness of the five point-tracking algorithm.

To substantially increase the robustness of the five point algorithm, a new set of parameters is created. These parameters take into consideration physical properties of the feature points. The property of symmetry is used to reduce the number of unknowns. Additionally, reasonable lower and upper bounds are placed on nose height and are represented as inequality constraints. As a result, the algorithm becomes more robust. Using these techniques, the number of unknowns is significantly reduced below 20.

Even though the following algorithm is described with respect to five feature points, it is straightforward to extend the idea to any number of feature points less than or greater than five feature points for improved robustness. Additionally, the algorithm can be applied to other objects besides a face as long as the other objects represent some level of symmetry. Head motion estimation is first described with respect to five feature points. Next, the algorithm is extended to incorporate other image point matches obtained from image registration methods.

Head Motion Estimation from Five Feature Points. Fig. 7 illustrates the new coordinate system used to represent feature points. E_1 202, E_2 204, M_1 206,

M₂ 208, and N 210 denote the left eye corner, right eye corner, left mouth corner, right mouth corner, and nose top, respectively. A new point E 212 denotes the midpoint between eye corners E₁, E₂ and a new point M 214 identifies the midpoint between mouth corners M₁, M₂. Notice that human faces exhibit some strong structural properties. For example, the left and right sides of a human face are very close to being symmetrical about the nose. Eye corners and mouth corners are almost coplanar. Based on these symmetrical characteristics, the following reasonable assumptions are made:

- (1) A line E₁E₂ connecting the eye corners E₁ and E₂ is parallel to a line M₁M₂ connecting the mouth corners.
- (2) A line centered on the nose (e.g., line EOM when viewed straight on or lines NM or NE when viewed from an angle as shown) is perpendicular to mouth line M₁M₂ and to eye line E₁E₂.

Let π be the plane defined by E₁, E₂, M₁ and M₂. Let O 216 denote the projection of point N on plane π . Let Ω_0 denote the coordinate system, which is originated at O with ON as the z-axis, OE as the y-axis; the x-axis is defined according to the right-hand system. In this coordinate system, based on the assumptions mentioned earlier, we can define the coordinates of E₁, E₂, M₁, M₂, N as $(-a, b, 0)^T$, $(a, b, 0)^T$, $(-d, -c, 0)^T$, $(d, -c, 0)^T$, $(0, 0, e)^T$, respectively.

By redefining the coordinate system, the number of parameters used to define five feature points is reduced from nine (9) parameters for generic five points to five (5) parameters for five feature points in this local coordinate system.

Let \mathbf{t} denote the coordinates of O under the camera coordinate system, and \mathbf{R} the rotation matrix whose three columns are vectors of the three coordinate axis of Ω_0 . For each point $\mathbf{p} \in \{E_1, E_2, M_1, M_2, N\}$, its coordinate under the camera

coordinate system is $\mathbf{R}\mathbf{p} + \mathbf{t}$. We call (\mathbf{R}, \mathbf{t}) the head pose transform. Given two images of the head under two different poses (assume the camera is static), let (\mathbf{R}, \mathbf{t}) and $(\mathbf{R}', \mathbf{t}')$ be their head pose transforms. For each point $\mathbf{p}_i \in \{E_1, E_2, M_1, M_2, N\}$, if we denote its image point in the first view by \mathbf{m}_i and that in the second view by \mathbf{m}'_i , we have the following equations:

$$\text{proj}(\mathbf{R}\mathbf{p}_i + \mathbf{t}) = \mathbf{m}_i \quad (1)$$

and

$$\text{proj}(\mathbf{R}'\mathbf{p}_i + \mathbf{t}') = \mathbf{m}'_i \quad (2)$$

where proj is the perspective projection. Notice that we can fix one of the coordinates a, b, c, d , since the scale of the head size cannot be determined from the images. As is well known, each pose has six (6) degrees of freedom. Therefore, the total number of unknowns is sixteen (16), and the total number of equations is 20. If we instead use their 3D coordinates as unknowns as in any typical bundle adjustment algorithms, we would end up with 20 unknowns and have the same number of equations. By using the generic properties of the face structure, the system becomes over-constrained, making the pose determination more robust.

To make the system even more robust, we add an inequality constraint on e . The idea is to force e to be positive and not too large compared to a, b, c, d . In the context of the face, the nose is always out of plane π . In particular, we use the following inequality:

$$0 \leq e \leq 3a \quad (3)$$

Three (3) is selected as the upper bound of e/a simply because it seems reasonable and it works well. The inequality constraint is finally converted to equality constraint by using a penalty function.

$$P_{nose} = \begin{cases} e * e & \text{if } e < 0 \\ 0 & \text{if } 0 \leq e \leq 3a \\ (e - 3a) * (e - 3a) & \text{if } e > 3a \end{cases} \quad (4)$$

In summary, based on equations (1), (2) and (4), we estimate a, b, c, d, e , (\mathbf{R}, \mathbf{t}) and $(\mathbf{R}', \mathbf{t}')$ by minimizing

$$F_{5pts} = \sum_{i=1}^5 w_i (\|\mathbf{m}_i - proj(\mathbf{R}\mathbf{p}_i + \mathbf{t})\|^2 + \|\mathbf{m}'_i - proj(\mathbf{R}'\mathbf{p}_i + \mathbf{t}')\|^2) + w_n P_{nose} \quad (5)$$

where w_i 's and w_n are the weighting factors, reflecting the contribution of each term. In our case, $w_i = 1$ except for the nose term which has a weight of 0.5 because it is usually more difficult to locate the nose top than other feature points. The weight for penalty w_n is set to 10. The objective function (5) is minimized using a Levenberg-Marquardt method [40]. More precisely, as mentioned earlier, we set a to a constant during minimization since the global head size cannot be determined from images.

Incorporating Image Point Matches. If we estimate camera motion using only the five user marked points, the result is sometimes not very accurate because the markers contain human errors. In this section, we describe how to incorporate the image point matches (obtained by any feature matching algorithm) to improve precision.

Let $(\mathbf{m}_j, \mathbf{m}'_j)$ ($j = 1 \dots K$) be the K point matches, each corresponding to the projections of a 3D point \mathbf{p}_j according to the perspective projection (1) and (2). 3D

points \mathbf{p}_j 's are unknown, so they are estimated. Assuming that each image point is extracted with the same accuracy, we can estimate $a, b, c, d, e, (\mathbf{R}, \mathbf{t}), (\mathbf{R}', \mathbf{t}')$, and $\{\mathbf{p}_j\} (j = 1 \dots K)$ by minimizing

$$F = F_{spts} + w_p \sum_{j=1}^K (\|\mathbf{m}_j - proj(\mathbf{R}\mathbf{p}_j + \mathbf{t})\|^2 + \|\mathbf{m}'_j - proj(\mathbf{R}'\mathbf{p}_j + \mathbf{t}')\|^2) \quad (6)$$

where F_{spts} is given by (5), and w_p is the weighting factor. We set $w_p = 1$ by assuming that the extracted points have the same accuracy as those of eye corners and mouth corners. The minimization can again be performed using a Levenberg-Marquardt method. This is a quite large minimization problem since we need to estimate $16 + 3K$ unknowns, and therefore it is computationally quite expensive especially for large K . Fortunately, as shown in [37], we can eliminate the 3D points using a first order approximation. The following term

$$\|\mathbf{m}_j - proj(\mathbf{R}\mathbf{p}_j + \mathbf{t})\|^2 + \|\mathbf{m}'_j - proj(\mathbf{R}'\mathbf{p}_j + \mathbf{t}')\|^2$$

can be shown to be equal, under the first order approximation, to

$$\frac{(\tilde{\mathbf{m}}_j^T \mathbf{E} \tilde{\mathbf{m}}_j)^2}{\tilde{\mathbf{m}}_j^T \mathbf{E}^T \mathbf{Z} \mathbf{Z}^T \mathbf{E} \tilde{\mathbf{m}}_j + \tilde{\mathbf{m}}_j^T \mathbf{E}^T \mathbf{Z} \mathbf{Z}^T \mathbf{E} \tilde{\mathbf{m}}_j}$$

where $\tilde{\mathbf{m}}_j = [\mathbf{m}_j^T, 1]^T$, $\tilde{\mathbf{m}}_j' = [\mathbf{m}_j'^T, 1]^T$, $\mathbf{Z} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{bmatrix}$, and \mathbf{E} is the essential matrix to be

defined below.

Let $(\mathbf{R}_r, \mathbf{t}_r)$ be the relative motion between two views. It is easy to see that

$$\mathbf{R}_r = \mathbf{R}'\mathbf{R}^T, \text{ and}$$

$$\mathbf{t}_r = \mathbf{t}' - \mathbf{R}'\mathbf{R}^T\mathbf{t}.$$

Furthermore, let's define a 3×3 antisymmetric matrix $[\mathbf{t}_r]_\times$ such that $[\mathbf{t}_r]_\times \mathbf{x} = \mathbf{t}_r \times \mathbf{x}$ for any 3D vector \mathbf{x} . The essential matrix is then given by

$$\mathbf{E} = [\mathbf{t}_r]_\times \mathbf{R}_r \quad (7)$$

which describes the epipolar geometry between two views [7].

In summary, the objective function (6) becomes

$$F = F_{pts} + w_p \sum_{j=1}^K \frac{(\tilde{\mathbf{m}}_j^T \mathbf{E} \tilde{\mathbf{m}}_j)^2}{\tilde{\mathbf{m}}_j^T \mathbf{E}^T \mathbf{Z} \mathbf{Z}^T \mathbf{E} \tilde{\mathbf{m}}_j + \tilde{\mathbf{m}}_j^T \mathbf{E}^T \mathbf{Z} \mathbf{Z}^T \mathbf{E} \tilde{\mathbf{m}}_j} \quad (8)$$

Notice that this is a much smaller minimization problem. We only need to estimate 16 parameters as in the five-point problem (5), instead of $16 + 3K$ unknowns.

To obtain a good initial estimate, we first use only the five feature points to estimate the head motion by using the algorithm described in Section 2. Thus we have the following two step algorithm:

Step1. Set $w_p=0$. Solve minimization problem 8.

Step2. Set $w_p=1$. Use the results of step1 as the initial estimates. Solve minimization problem (8).

Notice that we can apply this idea to the more general cases where the number of feature points is not five. For example, if there are only two eye corners and mouth corners, we'll end up with 14 unknowns and $16 + 3K$ equations. Other symmetric feature points (such as the outside eye corners, nostrils, and the like) can be added into equation 8 in a similar way by using the local coordinate system Ω_0 .

Head Motion Estimation Results

In this section, we show some test results to compare the new algorithm with the traditional algorithms. Since there are multiple traditional algorithms, we chose to implement the algorithm as described in [34]. It works by first computing an initial estimate of the head motion from the essential matrix [7], and then re-estimate the motion with a nonlinear least-squares technique.

We have run both the traditional algorithm and the new algorithm on many real examples. We found many cases where the traditional algorithm fails while the new algorithm successfully results in reasonable camera motions. When the traditional algorithm fails, the computed motion is completely bogus, and the 3D reconstructions give meaningless results. But the new algorithm gives a reasonable result. We generate 3D reconstructions based on the estimated motion, and perform Delauney triangulation.

We have also performed experiments on artificially generated data. We arbitrarily select 80 vertices from a 3D face model and project its vertices on two views (the head motion is eight degrees apart). The image size is 640 by 480 pixels. We also project the five 3D feature points (eye corners, nose top, and mouth corners) to generate the image coordinates of the markers. We then add random noises to the coordinates (u , v) of both the image points and the markers. The noises are generated by a pseudo-random generator subject to Gaussian distribution with zero mean and variance ranging from 0.4 to 1.2. We add noise to the marker's co-ordinates as well. The results are plotted in Figure 3. The blue curve shows the results of the traditional algorithm and the red curve shows the results of our new algorithm. The horizontal axis is the variance of the noise distribution.

1 The vertical axis is the difference between the estimated motion and the actual
2 motion. The translation vector of the estimated motion is scaled so that its
3 magnitude is the same as the actual motion. The difference between two rotations
4 is measured as the Euclidean distance between the two rotational matrices.

5 We can see that as the noise increases, the error of the traditional algorithm
6 has a sudden jump at certain point. But, the errors of our new algorithm grow
7 much more slowly.

8 9 **3D Reconstruction.**

10 Fig. 6 illustrates a step 114, where matched points are reconstructed in 3D
11 space with respect to the camera frame at the time when the first base image was
12 taken. Let $(\mathbf{m}, \mathbf{m}')$ be a couple of matched points, and \mathbf{p} be their corresponding
13 point in space. 3D point \mathbf{p} is estimated such that $\|\mathbf{m} - \hat{\mathbf{m}}\|^2 + \|\mathbf{m}' - \hat{\mathbf{m}}'\|^2$ is
14 minimized, where $\hat{\mathbf{m}}$ and $\hat{\mathbf{m}}'$ are projections of \mathbf{p} in both images according to the
15 equation $\lambda \tilde{\mathbf{m}} = \mathbf{A} \mathbf{P} \Omega \tilde{\mathbf{p}}$.

16 3D positions of the markers are determined in the same way.

17 18 **Fitting a Face Model**

19 This stage of processing creates a 3D model of the face. The face model
20 fitting process consists of two steps: fitting to 3D reconstructed points and fine
21 adjustment using image information.

22 23 **3D Fitting**

24 A step 120 comprises constructing a realistic 3D face model from the
25 reconstructed 3D image calculated in step 111. Given a set of reconstructed 3D

points from matched corners and markers, the fitting process applies a combination of deformation vectors to a pre-specified, neutral face model, to deform the neutral face model approximately to the reconstructed face model. The technique searches for both the *pose* of the face and the metric coefficients to minimize the distances from the reconstructed 3D points to the neutral face mesh. The pose of the face is the transformation $\mathbf{T} = \begin{pmatrix} s\mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{pmatrix}$ from the coordinate frame of the neutral face mesh to the camera frame, where \mathbf{R} is a 3×3 rotation matrix, \mathbf{t} is a translation, and s is a global scale. For any 3D vector \mathbf{p} , we use notation $\mathbf{T}(\mathbf{p}) = s\mathbf{R}\mathbf{p} + \mathbf{t}$.

The vertex coordinates of the face mesh in the camera frame is a function of both the metric coefficients and the pose of the face. Given metric coefficients (c_1, \dots, c_m) and pose \mathbf{T} , the face geometry in the camera frame is given by

$$\mathbf{S} = \mathbf{T} \left(\mathbf{S}^0 + \sum_{i=1}^n c_i \mathbf{M}^i \right)$$

Since the face mesh is a triangular mesh, any point on a triangle is a linear combination of the three triangle vertexes in terms of barycentric coordinates. So any point on a triangle is also a function of \mathbf{T} and metric coefficients. Furthermore, when \mathbf{T} is fixed, it is simply a linear function of the metric coefficients.

Let $(\mathbf{p}_1, \mathbf{p}_2, \dots, \mathbf{p}_k)$ be the reconstructed corner points, and $(\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_5)$ be the reconstructed markers. Denote the distance from \mathbf{p}_i to the face mesh \mathbf{S} by $d(\mathbf{p}_i, \mathbf{S})$. Assume marker \mathbf{q}_j corresponds to vertex \mathbf{v}_{m_j} of the face mesh, and denote the distance between \mathbf{q}_j and \mathbf{v}_{m_j} by $d(\mathbf{q}_j, \mathbf{v}_{m_j})$. The fitting process consists of finding pose \mathbf{T} and metric coefficients $\{c_1, \dots, c_n\}$ by minimizing

$$\sum_{i=1}^n w_i d^2(\mathbf{p}_i, \mathbf{S}) + \sum_{j=1}^5 d^2(\mathbf{q}_j, \mathbf{v}_{m_j})$$

where w_i is a weighting factor.

1 To solve this problem, we use an iterative closest point approach. At each
2 iteration, we first fix \mathbf{T} . For each \mathbf{p}_i , we find the closest point \mathbf{g}_i on the current face
3 mesh \mathbf{S} . We then minimize $\sum w_i d^2(p_i, S) + \sum d^2(\mathbf{q}_j, \mathbf{v}_{m_j})$. We set w_i to be 1 at
4 the first iteration and $1.0/1+d^2(\mathbf{p}_i, \mathbf{g}_i)$ in the subsequent iterations. The reason for
5 using weights is that the reconstruction from images is noisy and such a weight
6 scheme is an effective way to avoid overfitting to the noisy data [8]. Since both \mathbf{g}_i
7 and \mathbf{v}_{m_j} are linear functions of the metric coefficients for fixed \mathbf{T} , the above
8 problem is a linear least square problem. We then fix the metric coefficients, and
9 solve for the pose. To do that, we recompute \mathbf{g}_i using the new metric coefficients.
10 Given a set of 3D corresponding points $(\mathbf{p}_i, \mathbf{g}_i)$ and $(\mathbf{q}_j, \mathbf{v}_{m_j})$, there are well known
11 algorithms to solve for the pose. We use the quaternion-based technique described
12 in [11]. To initialize this iterative process, we first use the 5 markers to compute
13 an initial estimate of the pose. In addition, to get a reasonable estimate of the head
14 size, we solve for the head-size related metric coefficients such that the resulting
15 face mesh matches the bounding box of the reconstructed 3D points.
16 Occasionally, the corner matching algorithm may produce points not on the face.
17 In that case, the metric coefficients will be out of the valid ranges, and we throw
18 away the point that is the most distant from the center of the face. We repeat this
19 process until metric coefficients become valid.

21 Fine Adjustment Using Image Information

22 After the geometric fitting process, we have now a face mesh that is a close
23 approximation to the real face. To further improve the result, we perform a search
24 for silhouettes and other face features in the images and use them to refine the
25 face geometry. The general problem of locating silhouettes and face features in

1 images is difficult, and is still a very active research area in computer vision.
2 However, the face mesh that we have obtained provides a good estimate of the
3 locations of the face features, so we only need to perform search in a small region.

4 We use the snake approach [15] to computer the silhouettes of the face. The
5 silhouette of the current face mesh is used as the initial estimate. For each point on
6 this piecewise linear curve, we find the maximum gradient location along the
7 normal direction within a small range (10 pixels each side in our implementation).
8 Then we solve for the vertexes (acting as control points) to minimize the total
9 distance between all the points and their corresponding maximum gradient
10 locations.

11 We use a similar approach to find the upper lips.

12 To find the outer eye corner (not marked), we rotate the current estimate of
13 that eye corner (given by the face mesh) around the marked eye corner by a small
14 angle, and look for the eye boundary using image gradient information. This is
15 repeated for several angles, and the boundary point that is the most distant to the
16 marked corner is chosen as the outer eye corner.

17 We could also use the snake approach to search for eyebrows. However,
18 our current implementation uses a slightly different approach. Instead of
19 maximizing image gradients across contours, we minimize the average intensity of
20 the image area that is covered by the eyebrow triangles. Again, the vertices of the
21 eyebrows are only allowed to move in a small region bounded by their neighboring
22 vertices. This has worked very robustly in our experiments.

23 We then use the face features and the image silhouettes as constraints in our
24 system to further improve the mesh, in a step 131. Notice that each vertex on the
25 mesh silhouette corresponds to a vertex on the image silhouette. We cast a ray

1 from the camera center through the vertex on the image silhouette. The projection
2 of the corresponding mesh vertex on this ray acts as the target position of the mesh
3 vertex. Let \mathbf{v} be the mesh vertex and \mathbf{h} the projection. We have equation $\mathbf{v} = \mathbf{h}$.
4 For each face feature, we obtain an equation in a similar way. These equations are
5 added to equation (5). The total set of equations is solved as before, i.e., we first
6 fix the post \mathbf{T} and use a linear least square approach to solve the metric
7 coefficients, and then fix the metric coefficients while solving for the pose.

9 Face Texture From Video Sequence

10 Now we have the geometry of the face from only two views that are close to
11 the frontal position. For the sides of the face, the texture from the two images is
12 therefore quite poor or even not available at all. Since each image only covers a
13 portion of the face, we need to combine all the images in the video sequence to
14 obtain a complete texture map. This is done by first determining the head pose for
15 the images in the video sequence and then blending them to create a complete
16 texture map.

18 Determining Head Motions in Video Sequences

19 Fig. 6 shows operations in creating a texture map. In an operation 140,
20 successive images are first matched using the same corner detection, corner
21 matching, and false match detection techniques described above. We could
22 combine the resulting motions incrementally to determine the head pose.
23 However, this estimation is quite noisy because it is computed only from 2D
24 points. As we already have the 3D face geometry, a more reliable pose estimation
25 can be obtained by combining both 3D and 2D information, as follows.

In an operation 141, the pose of each successive image is determined. Let us denote the first base image by I_0 . This base image comprises one of the two initial still images, for which the pose is already known. Because we know the pose of the base image, we can determine the 3D position of each point in the base image relative to the facial model that has already been computed.

We will denote the images on the video sequences by I_1, \dots, I_v . The relative head motion from I_{i-1} to I_i is given by $R = \begin{pmatrix} \mathbf{R}_{ri} & \mathbf{t}_{ri} \\ \mathbf{0}^T & 1 \end{pmatrix}$, and the head pose corresponding to image I_i with respect to the camera frame is denoted by Ω_i . The technique works incrementally, starting with I_0 and I_1 . For each pair of images (I_{i-1}, I_i), we perform a matching operation to match points of image I_i with corresponding points in I_{i-1} . This operation uses the corner matching algorithm described above. We then perform a minimization operation, which calculates the pose of I_i such that projections of 3D positions of the matched points of I_{i-1} onto I_i coincide approximately with the corresponding matched points of I_i . More specifically, the minimization operation minimizes differences between the projections of 3D positions of the matched points of I_{i-1} onto I_i and the corresponding matched points of I_i . Let us denote the matched corner pairs as $\{(m_j, m'_j) \mid j = 1, \dots, l\}$. For each \mathbf{m}_j in I_{i-1} , we cast a ray from the camera center through \mathbf{m}_j , and compute the intersection \mathbf{x}_j of that ray with the face mesh corresponding to image I_{i-1} . According to the equation $\lambda \tilde{\mathbf{m}} = \mathbf{A} \mathbf{P} \Omega \tilde{\mathbf{p}}$, R_i is subject to the following equations

$$\mathbf{A} \mathbf{P} R_i \tilde{\mathbf{x}}_j = \lambda_j \tilde{\mathbf{m}}'_j \quad \text{for } j = 1, \dots, l$$

where \mathbf{A} , \mathbf{P} , \mathbf{x}_j and \mathbf{m}'_j are known. Each of the above equations gives two constraints on R_i . We compute R_i with a technique described in [37], which minimizes the sum of differences between each pair of matched points (m_j, m'_j) .

After R_i is computed, the head pose for image I_i in the camera frame is given by $\Omega_i = R_i \Omega_{i-1}$. The head pose Ω_0 is known from previous calculations involving the two still images.

In general, it is inefficient to use all the images in the video sequence for texture blending, because head motion between two consecutive frames is usually very small. To avoid unnecessary computation, the following process is used to automatically select images from the video sequence. Let us call the amount of rotation of the head between two consecutive frames the *rotation speed*. If s is the current rotation speed and α is the desired angle between each pair of selected images, the next image is selected α/s frames away. In our implementation, the initial guess of the rotation speed is set to 1 degree/frame and the desired separation angle is equal to 5 degrees.

Texture Blending

Operation 142 is a texture blending operation. After the head pose of an image is computed, we use an approach similar to Pighin et al.'s method [26] to generate a view independent texture map. We also construct the texture map on a virtual cylinder enclosing the face model. But instead of casting a ray from each pixel to the face mesh and computing the texture blending weights on a pixel by pixel basis, we use a more efficient approach. For each vertex on the face mesh, we computed the blending weight for each image based on the angle between surface normal and the camera direction [26]. If the vertex is invisible, its weight is set to 0.0. The weights are then normalized so that the sum of the weights over all the images is equal to 1.0. We then set the colors of the vertexes to be their weights, and use the rendered image of the cylindrical mapped mesh as the weight

1 map. For each image, we also generate a cylindrical texture map by rendering the
2 cylindrical mapped mesh with the current image as texture map. Let C_i and
3 W_i ($i = 1, \dots, k$) be the cylindrical texture maps and the weight maps. Let D be
4 the final blended texture map. For each pixel (u, v) , its color on the final blended
5 texture map is

$$C(u, v) = \sum_{i=1}^k W_i(u, v) C_i(u, v).$$

7 Because the rendering operations can be done using graphics hardware, this
8 approach is very fast.

10 User Interface

11 We have built a user interface to guide the user through collecting the
12 required images and video sequences, and marking two images. The generic head
13 model without texture is used as a guide. Recorded instructions are lip-synced
14 with the head directing the user to first look at a dot on the screen and push a key
15 to take a picture. A second dot appears and the user is asked to take the second
16 still image. The synthetic face mimics the actions the user is to follow. After the
17 two still images are taken, the guide directs the user to slowly turn his/her head to
18 record the video sequences. Finally, the guide places red dots on her own face and
19 directs the user to do the same on the two still images. The collected images and
20 markings are then processed and a minute or two later they have a synthetic head
21 that resembles them.

Animation

Having obtained the 3D textured face model, the user can immediately animate the model with the application of facial expressions including frowns, smiles, mouth open, etc.

To accomplish this we have defined a set of vectors, which we call *posemes*. Like the metric vectors described previously, posemes are a collection of artist-designed displacements. We can apply these displacements to any face as long as it has the same topology as the neutral face. Posemes are collected in a library of actions and expressions.

The idle motions of the head and eyeballs are generated using Perlin's noise functions [24, 25].

Conclusions

We have developed a system to construct textured 3D face models from video sequences with minimal user intervention. A new head motion estimation algorithm takes advantage of the physical properties of human face features. The algorithm significantly improves the robustness over traditional motion estimation methodologies. It can be applied to human face modeling and tracking systems where the markers can be obtained either through user intervention or by using automatic feature detection algorithms. This algorithm can be easily extended to general cases where the number of feature points is not necessarily five.

Although details of specific implementations and embodiments are described above, such details are intended to satisfy statutory disclosure obligations rather than to limit the scope of the following claims. Thus, the invention as defined by the claims is not limited to the specific features described

1 above. Rather, the invention is claimed in any of its forms or modifications that
2 fall within the proper scope of the appended claims, appropriately interpreted in
3 accordance with the doctrine of equivalents.

4 **REFERENCES**

- 5
- 6 [1] T. Akimoto, Y. Suenaga, and R. S. Wallace. Automatic 3d facial models.
7 *IEEE Computer Graphics and Applications*, 13(5):16-22, September 1993.
- 8 [2] J. Barron, D. Fleet, and S. Beauchemin. Performance of optical flow
9 techniques. *The International Journal of Computer Vision*, 12(1):43-77,
10 1994.
- 11 [3] V. Blanz and T. Vetter. A morphable model for the synthesis of 3d faces.
12 In *Computer Graphics, Annual Conference Series*, pages 187-194.
13 Siggraph, August 1999
- 14 [4] B. Dariush, S. B. Kang, and K. Waters. Spatiotemporal analysis of face
15 profiles: Detection, segmentation, and registration. In *Proc. of the 3rd*
16 *International Conference on Automatic Face and Gesture Recognition*,
17 pages 248-253. IEEE, April 1998.
- 18 [5] D. DeCarlo, D. Metaxas, and M. Stone. An anthropometric face model
19 using variational techniques. In *Computer Graphics, Annual Conference*
20 *Series*, pages 67-74. Siggraph, July 1998.
- 21 [6] S. DiPaola. Extending the range of facial types. *Journal of Visualization*
22 *and Computer Animation*, 2(4):129-131, 1991.
- 23
24
25

- 1 [7] O. Faugeras. *Three-Dimensional Computer Vision: a Geometric Viewpoint*.
2 MIT Press, 1993.
- 3 [8] P. Fua and C. Miccio. From regular images to animated heads: A least
4 squares approach. In *Eurographics of Computer Vision*, pages 188-202,
5 1996.
- 6 [9] B. Guenter, C. Grimm, D. Wood, H. Malvar, and F. Pighin. Making faces.
7 In *Computer Graphics, Annual Conference Series*, pages 55-66. Siggraph,
8 July 1998.
- 9 [10] C. Harris and M. Stephens. A combined corner and edge detector. In *Proc.*
10 *4th Alvey Vision Conf.*, pages 189-192, 1988.
- 11 [11] B. K. Horn. Closed-form Solution of Absolute Orientation using Unit
12 Quaternions. *Journal of Optical Society A*, 4(4):629-642, Apr. 1987.
- 13 [12] B. K. P. Horn and B. G. Schunk. Determining Optical Flow. *Artificial*
14 *Intelligence*, 17:185-203, 1981.
- 15 [13] H. H. S. Ip and L. Yin. Constructing a 3d individualized head model from
16 two orthogonal views. *The Visual Computer*, (12):254-266, 1996.
- 17 [14] S. B. Kang and M. Jones. Appearance-based structure from motion using
18 linear classes of 3-d models. *Manuscript*, 1999.
- 19 [15] M. Kass, A. Witkin, and D. Terzopoulos. SNAKES: Active contour
20 models. *The International Journal of Computer Vision*, 1:321-332, Jan.
21 1988.
- 22
23
24
25

- 1 [16] A. Lanitis, C. J. Taylor, and T. F. Cootes. Automatic interpretation and
2 coding of face images using flexible models. *IEEE Transactions on*
3 *Pattern Analysis and Machine Intelligence*, 19(7):743-756, 1997.
- 4 [17] Y. C. Lee, D. Terzopoulos, and K. Waters. Constructing physics-based
5 facial models of individuals. In *Proceedings of Graphics Interface*, Pages
6 1-8, 1993.
- 7 [18] Y. C. Lee, D. Terzopoulos, and K. Waters. Realistic modeling for facial
8 animation. In *Computer Graphics, Annual Conference Series*, pages 55-62.
9 SIGGRAPH, 1995.
- 10 [19] J. P. Lewis. Algorithms for solid noise synthesis. In *Computer Graphics,*
11 *Annual Conference Series*, pages 263-270. Siggraph, 1989.
- 12 [20] N. Magneneat-Thalmann, H. Minh, M. Angelis, and D. Thalmann.
13 Design, transformation and animation of human faces. *Visual Computer*, (5):32-
14 39, 1989.
- 15 [21] F. I. Parke. Computer generated animation of faces. In *ACM National*
16 *Conference*, November 1972
- 17 [22] F. I. Parke. *A Parametric Model of human Faces*. PhD thesis, University
18 of Utah, 1974
- 19 [23] F. I. Parke and K. Waters. *Computer Facial Animation*. AKPeters,
20 Wellesley, Massachusetts, 1996.
- 21
22
23
24
25

- 1 [24] K. Perlin. Real time responsive animation with personality. *IEEE*
2 *Transactions on Visualization and Computer Graphics*, 1(1), 1995.
- 3 [25] K. Perlin and A. Goldberg. Improv: A system for scripting interactive
4 actors in virtual worlds. In *Computer Graphics, Annual Conference Series*,
5 pages 205-216. Siggraph, August 1995.
- 6 [26] F. Pighin, J. Hecker, D. Lischinski, R. Szeliski, and D.H. Salesin.
7 Synthesizing realistic facial expressions from photographs. In *Computer*
8 *Graphics, Annual Conference Series*, pages 75-84. Siggraph, July 1998.
- 9 [27] S. Platt and N. Badler. Animating facial expression. *Computer*
10 *Graphics*, 15(3):245-252, 1981.
- 11 [28] P. Rousseeuw and A. Leroy. *Robust Regression and Outlier Detection*.
12 John Wiley & Sons, New York, 1987.
- 13 [29] T. Shakunaga, K. Ogawa, and S. Oki. Integration of eigentemplate and
14 structure matching for automatic facial feature detection. In *Proc. of the 3rd*
15 *International Conference on Automatic Face and Gesture Recognition*,
16 pages 94-99, April 1998.
- 17 [30] D. Terzopoulos and K. Waters. Physically based facial modeling, analysis,
18 and animation. In *Visualization and Computer Animation*, pages 73-80,
19 1990.
- 20 [31] J. T. Todd, S. M. Leonard, R. E. Shaw, and J. B. Pittenger. The perception
21 of human growth. *Scientific American*, (1242):106-114, 1980.
- 22
23
24
25

- 1 [32] T. Vetter and T. Poggio. Linear object classes and image synthesis from a
2 single example image. *IEEE Transactions on Pattern Analysis and*
3 *Machine Intelligence*, 19(7):733-742, 1997.
- 4 [33] K. Waters. A muscle model for animating three-dimensional facial
5 expression. *Computer Graphics*, 22(4):17-24, 1987.
- 6 [34] Z. Zhang. Motion and structure from two perspective views: From
7 essential parameters to euclidean motion via fundamental matrix. *Journal*
8 *of the Optical Society of America A*, 14(11):2938-2950, 1997.
- 9 [35] Z. Zhang. Determining the epipolar geometry and its uncertainty: A review.
10 *The International Journal of Computer Vision*, 27(2):161-195, 1998.
- 11 [36] Z. Zhang. Flexible camera calibration by viewing a plane from unknown
12 orientations. In *International Conference on Computer Vision (ICCV'99)*,
13 pages 666-673, 1999.
- 14 [37] J. Y. Zheng. Acquiring 3-d models from sequences of contours. *IEEE*
15 *Transactions of Pattern Analysis and Machine Intelligence*, 16(2):163-178,
16 February 1994.
- 17 [38] P. Fua. Using model-driven bundle-adjustment to model heads from raw
18 video sequences. In *International Conference on Computer Vision*, pages
19 46-53, Sept. 1999.
- 20 [39] T. S. Jebara and A. Pentland. Parameterized structure from motion for 3d
21 adaptive feedback tracking of faces. In *Proc. CVPR*, pages 144-150, 1997.
- 22
23
24
25

- 1 [40] J. More. The levenberg-marquardt algorithm, implementation and theory. In
2 G. A. Watson, editor, Numerical Analysis, Lecture Notes in Mathematics
3 630. Springer-Verlag, 1977.
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25